



The EPIC Community Workshop

Tim Carroll



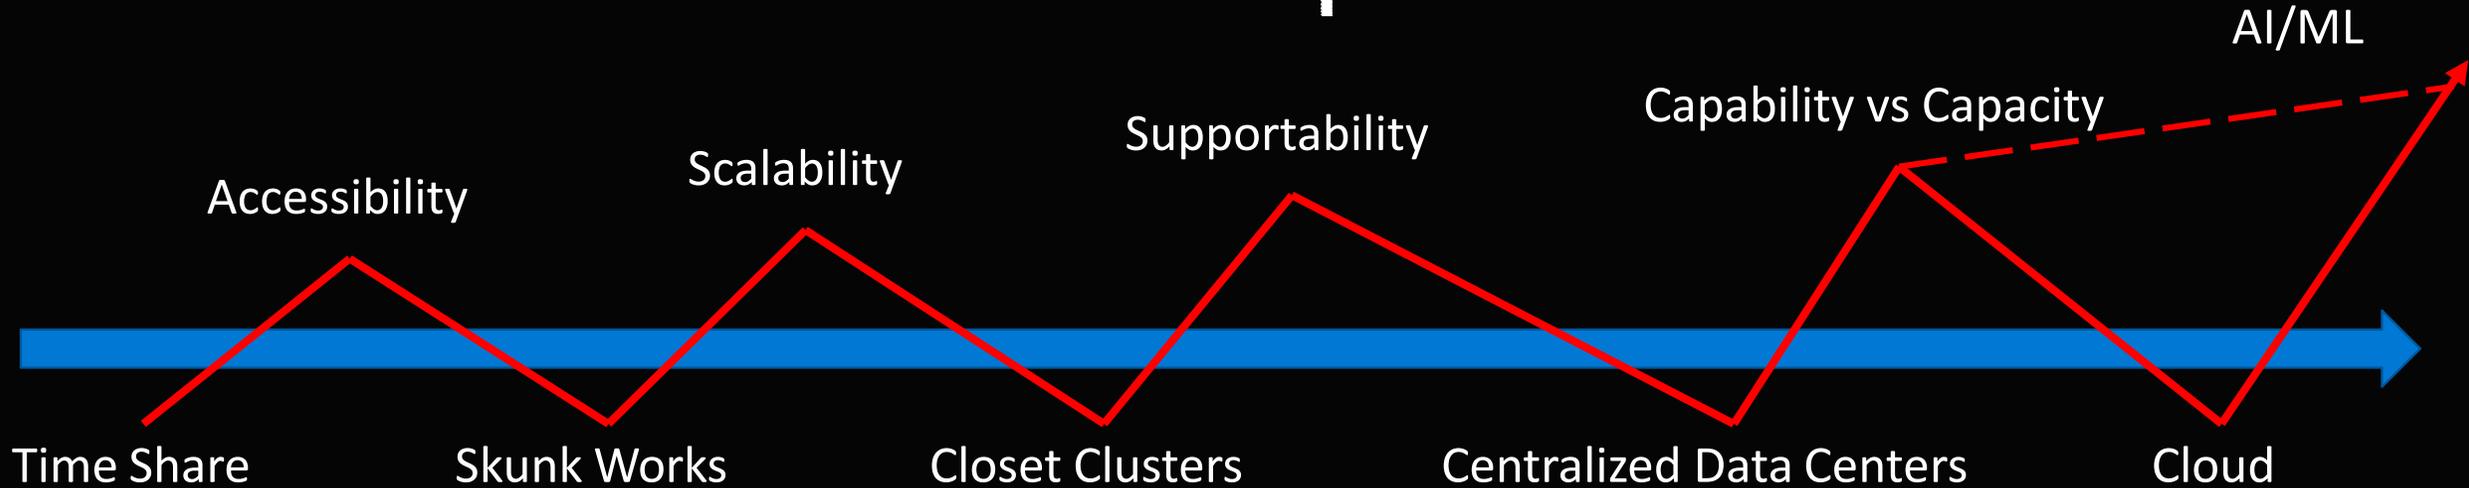


“Perhaps some day in the dim future, it will be possible to advance the computations faster than the weather advances and at a cost less than the saving to mankind due to the information gained...but that is a dream”

- Lewis Fry Richardson, 1922

Excerpted from “The Weather Machine” by Andrew Blum

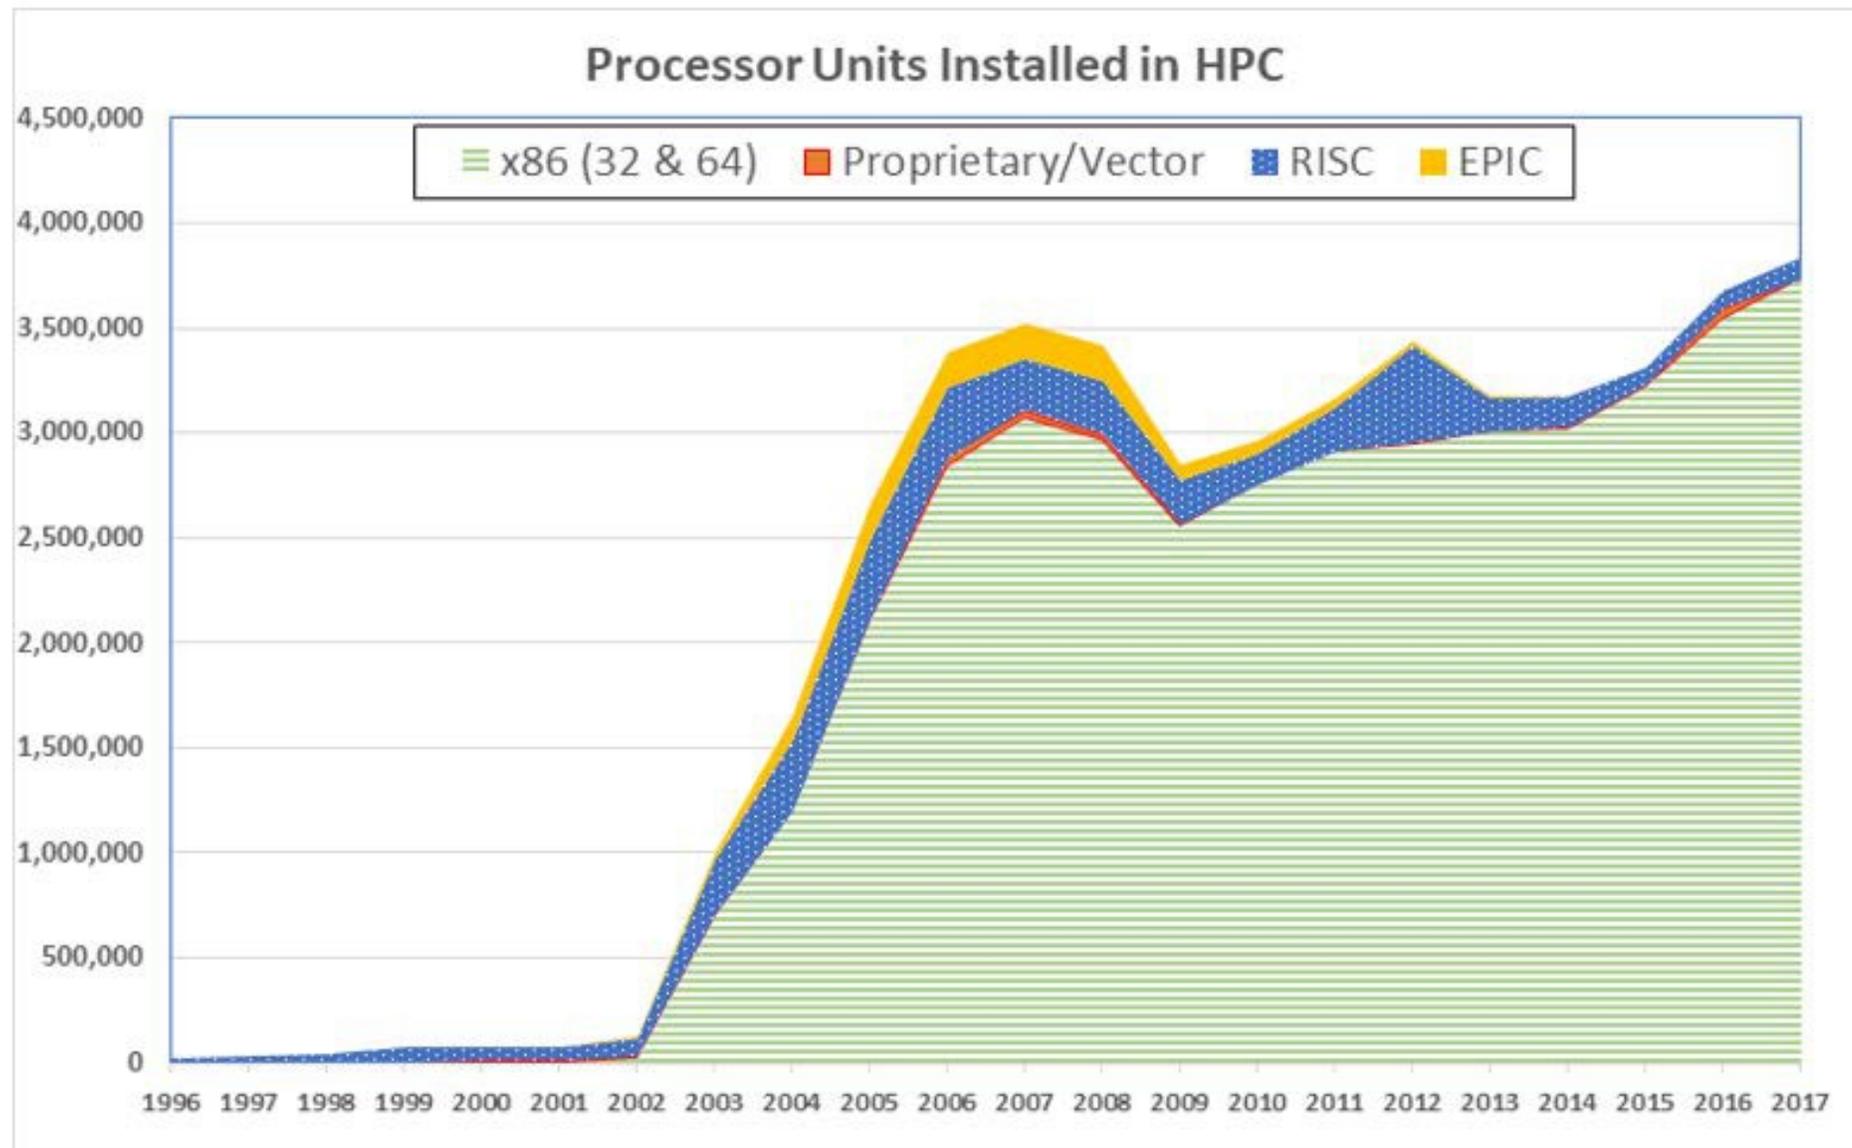
User Problems Drive Roadmap



- HPC resources evolve only when a barrier to answers needs to be removed
- Users not vendors drive the solution
- Cloud initially tried to serve a user base who was already well served
 - Oversubscribed web servers are not a barrier to answers
- Centralized Data centers replaced closet clusters; cloud will not replace data centers

**The Tipping
Point was
Access
Not Price,
Driven by
Community
Not Vendors**

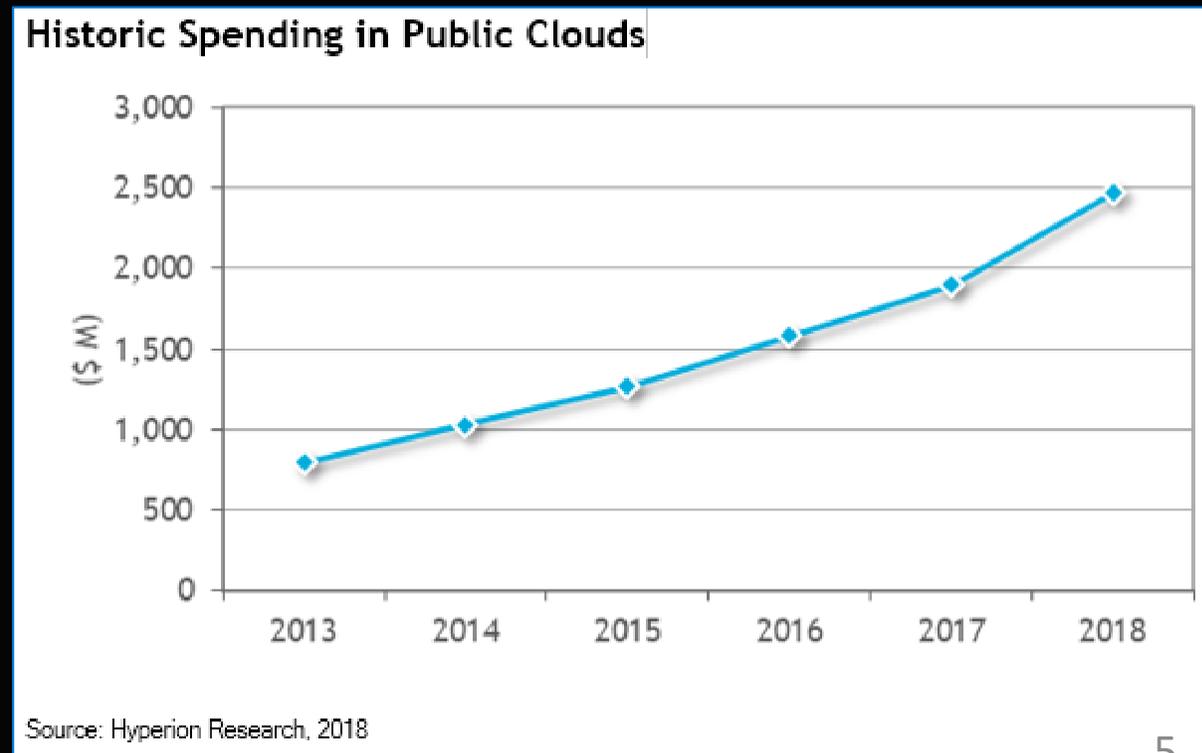
Processor Units Installed in HPC from 1996 to 2017



Source: Hyperion Research, 2018

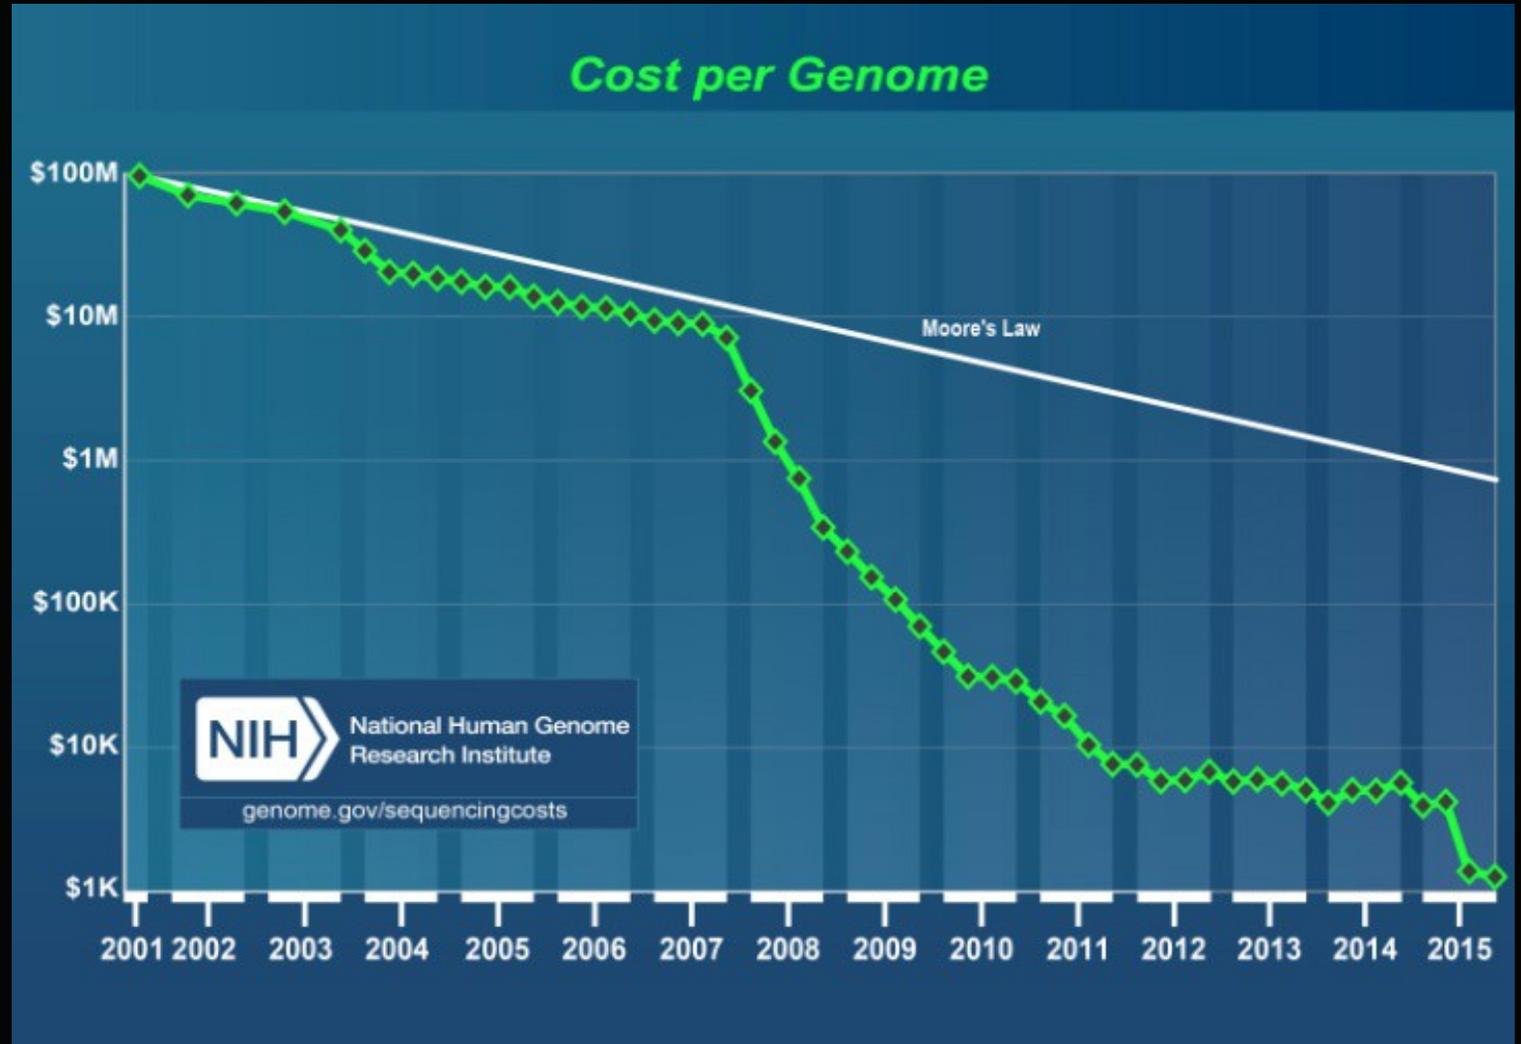
- Over 70% of HPC sites run some jobs in public clouds
 - Up from 13% in 2011
- Just over 10% of all HPC jobs are now running in clouds
 - Public clouds are cost effective for some jobs, but up to 10x more expensive for others
 - Key concerns: security, data loss

Hyperion Observations: HPC in the Cloud



Commodity clusters fueled genomics.

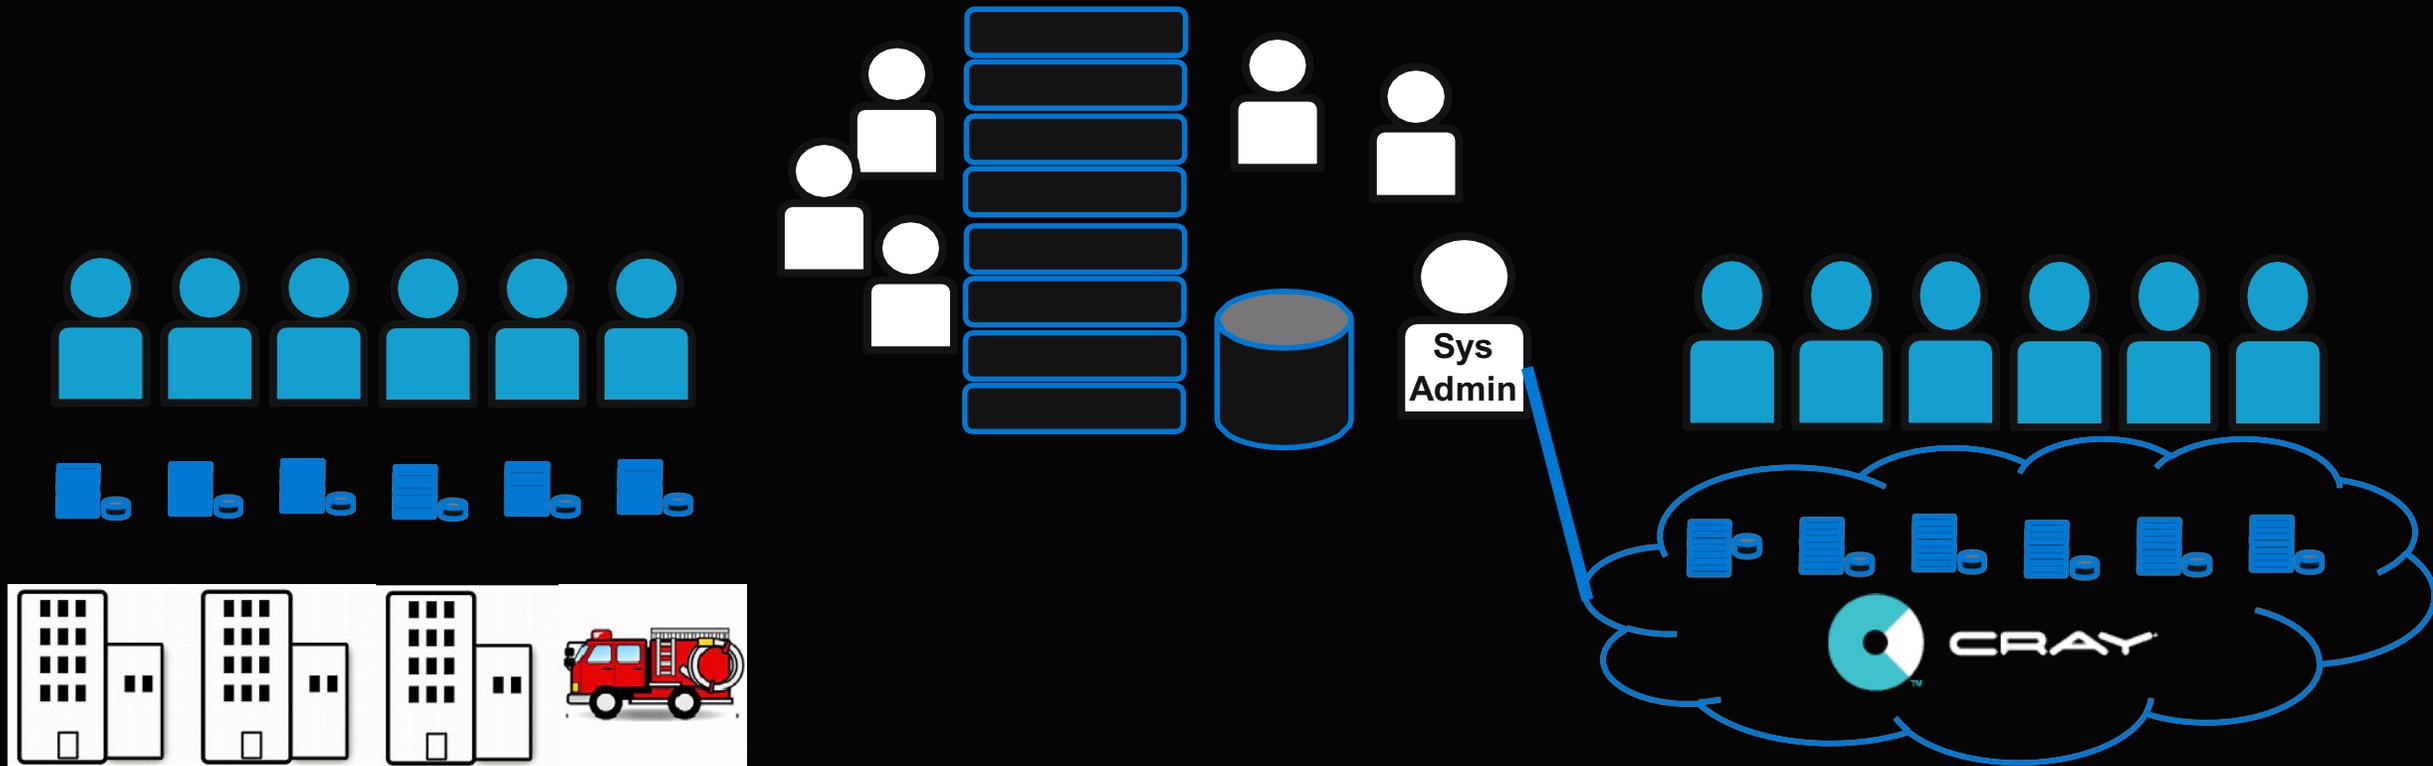
Cloud can fuel weather.



Closet Clusters Weren't All Bad

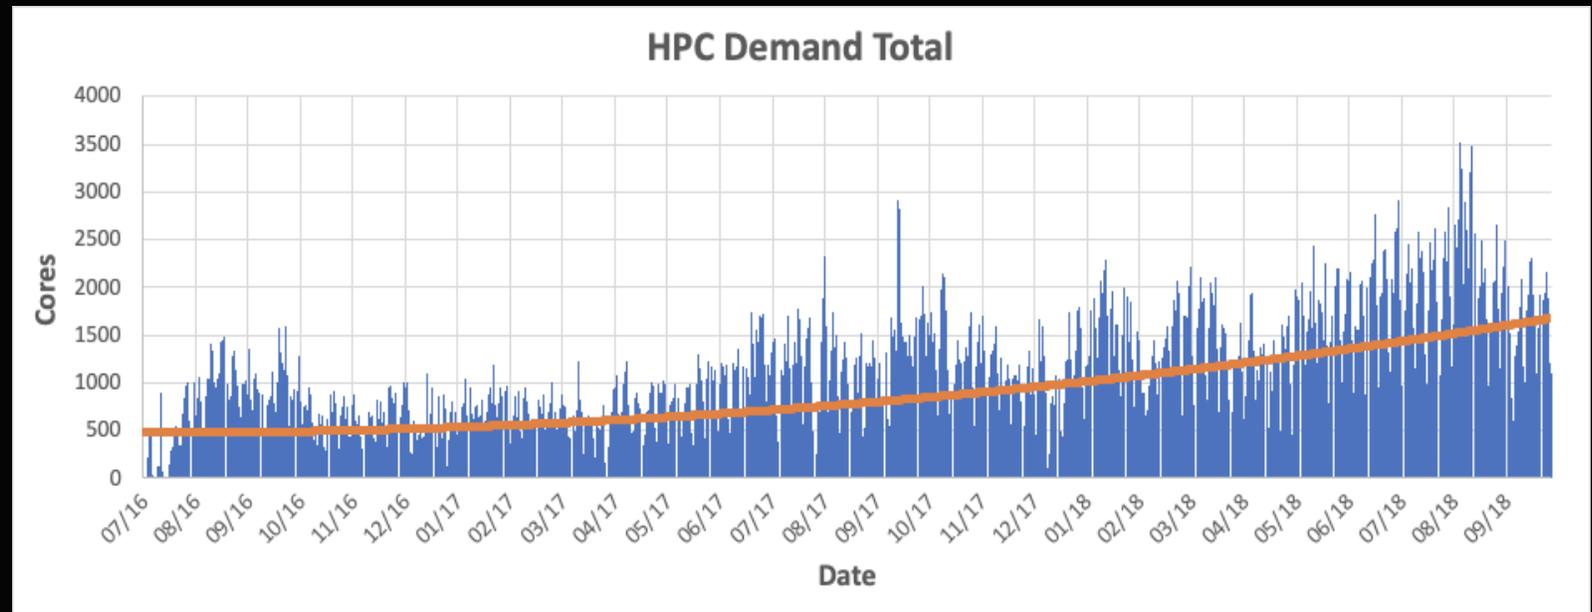
Workload level control for users

System level control for administration



Costing and Budgeting

- Costing is the end, not the beginning
- Assess vertically, not horizontally
 - Start at workflows and work back
 - Real costs are de risk traditional extrapolation methods
- Greater cost savings than negotiating



Microsoft to invest \$1B in OpenAI



“We want AGI to work with people to solve currently intractable multidisciplinary problems, including global challenges such as climate change, affordable and high-quality health care, and personalized education,” – *OpenAI blog post*

The two companies will work to build out Microsoft’s cloud-computing platform, Azure, to create supercomputing technologies strong enough to support the kind of innovations in artificial intelligence that might one day lead to AGI. Limitations in computing power are a major barrier to AI development – *Wired Magazine*

Thank you

Tim.Carroll@Microsoft.com

Review Questions

1. What is the most important action we can undertake to prepare our legacy codes for execution in the cloud?
2. In your view, how can NOAA most efficiently and effectively utilize cloud HPC? Status of their UFS implementations in the cloud. What will be the extent of their user support? What is their HPC capacity?
3. What new technologies do you see enhancing full Earth system modeling?
4. How should the cloud intersect with local HPC resources?
5. Will the cloud ever be seen as a platform for operational models?
6. What are impediments to adoption of new technologies?
7. Describe your vision for implementing a cloud solution for both NOAA & its community partners through EPIC.
8. What cloud strategies do you see being most successful within a private public partnership?

1. Collaborate with the vendor community
2. Inventory and catalogue all workloads currently running in operational and R&D
3. The user community will decide which technologies will enhance – *could be GitHub*
4. Cloud should be a local HPC resource
5. Yes
6. Cloud is not a technology; it is a service delivery model. Don't use the traditional HPC procurement model to procure cloud services.
7. What is problem to be solved?
8. What is problem to be solved?